GPCG
General Practice Computing Group

1st edition June 2005

www.gpcg.org

Medical Practice Network Security
Firewall Tutorial

Medical Practice Network Security - Firewall Tutorial

INTRODUCTION **4**

About this tutorial 4

What are firewalls? 4

Why do you need them? 4

What other computer security do you need? 5

What if your ISP already provides a firewall? 6

Firewall implementation issues 7

Do you have the necessary IT skills in-house? 7

STEP 1: Understanding firewalls in principle **8**

STEP 2: Understanding how the Internet works **9**

STEP 3: Deciding which firewall product you need **11**

STEP 4: Understanding firewall technologies **14**

STEP 5: Understanding different types of firewalls **15**

STEP 6: Understanding network addressing **18**

STEP 7: Understanding ports and firewall configuration **20**

STEP 8: Suggested firewall products **22**

STEP 9: Principles of firewall configuration **25**

STEP 10: DIY security audits **26**

STEP 11: How to audit your firewall - step by step **30**

STEP 12: Firewall checklists - after installation **37**

FURTHER INFORMATION **38**

Virtual Private Network 38

Failover/load balancing 40

GLOSSARY **42**

Acknowledgements

The General Practice Computing Group would like to thank the following people for contributing to *Medical Practice Network Security – Firewall Tutorial*. This resource has been developed as supporting information to the *GPCG Computer Security – Firewall Guideline,* a companion document to the *GPCG Computer Security Self-Assessment Guideline and Checklist for General Practitioners* (the Security Guidelines). Dr Horst Herb formulated the original LAN Firewalls document with subsequent input from Dr Ian Cheong, Dr Rob Hosking and Dr David Guest. Further technical expertise was received from the Broadband for Health Section of the Department of Health and Ageing. Additional feedback has been provided by state-based officers of the Australian Divisions of General Practice.

Medical Practice Network Security – Firewall Tutorial was jointly funded by the Australian Government and General Practice Computing Group.

General Practice Computing Group
C/- Royal Australian College of General Practitioners
1 Palmerston Crescent
South Melbourne, Vic 3205
Tel: (03) 8699 0414

www.gpcg.org.au

© June 2005

INTRODUCTION

About this tutorial

The information in this tutorial has been put together by the General Practice Computing Group (GPCG) with additional input provided by the Broadband for Health section of the Australian Department of Health and Ageing and State-based officers of the Australian Divisions of General Practice.

It is a reference for practice managers, IT service providers and GPs to help you:

- understand more about firewalls and why we need them.
- select, install, configure and maintain the firewall best suited to your medical practice.

While this tutorial can enhance awareness about firewalls and the need for them, you will still require the appropriate technical expertise to follow through and properly protect your computer system.

What is a firewall?

A firewall is a system designed to prevent unauthorised access to or from a private network (e.g. between your practice network and the Internet). Firewalls can be implemented in both hardware and software, or a combination of the two. All messages entering or leaving the private network must pass through the firewall, which examines each message and blocks those that do not meet the specified security criteria.

There are several types of firewall techniques:

- **Packet filter** – a packet filter examines each packet (message) entering or leaving the network and accepts or rejects it based on the packet type or source/destination address, according to user-defined rules. Packet filtering is fairly effective and transparent to users, but it is difficult to configure. In addition, it is susceptible to IP spoofing, i.e. using a legitimate IP address or packet to gain unauthorised access to a computer.
- **Stateful inspection** – a stateful inspector monitors the state of network connections that pass through the firewall. It inspects incoming and outgoing packets to determine if they correspond to an authorised connection.
- **Application proxy** – an application proxy only permits packets related to specific applications to pass through the firewall. For example, SMTP packets for email and HTTP for web browsing. This is very effective, but can reduce the computer's performance.

In practice, many firewalls use two or all of these techniques.

Why do you need them?

It's the law

Australian privacy legislation requires medical practices take reasonable steps to protect confidential patient data. If your practice computer system/s connects to the Internet, the GPCG recommends you protect that connection with a properly configured firewall.

Security breaches cost you

According to the 2004 Australian Computer Crime and Security Survey by the Australian Computer Emergency Response Team (AUSCERT), the key computer security trends in Australia are:

- 95% of respondents reported experiencing computer security incidents in the past 12 months, with the majority of organisations experiencing between one and five incidents.
- the number of respondents experiencing attacks that harmed the confidentiality, integrity and availability of networks, data or systems increased from 42% in 2003 to 49% in 2004.
- the average financial loss per incident was $116 212.
- 88% of attacks originate from external sources.
- 13% of respondents reported that hackers had penetrated their systems.

The full survey results are at www.auscert.org.au

Security breaches compromise your practice

It is easy for intruders to gain full control over unprotected computers connected to the Internet. An intruder's primary goal may not be the data you have stored on your computer. They may want to use your system to deliver spam mail for example. However, your data is still at risk of being compromised.

Most computers in medical practice store confidential patient data. Even if a compromised computer doesn't store confidential data, it can potentially be used to access the rest of the practice network or even other health networks that your practice connects to.

If a computer is connected to the Internet, even only temporarily, a firewall is essential. It is possible for an attack to occur in very little time and at any hour of the day. Such attacks may not be obvious to the user.

What other computer security do you need?

Firewalls don’t stop all threats

One of the most reputable network security institutions, SANS (the SysAdmin, Audit, Network, Security Institute), regularly publishes lists of the worst mistakes people make that lead to security breaches – see www.sans.org/resources/mistakes.php.

According to SANS, one of the worst mistakes made by Senior Executives that leads to security breaches is to rely primarily on a firewall. A firewall is considered a first line of defence in protecting private information against network attacks. While firewalls can be configured to filter certain types of traffic, which does not necessarily mean they stop all threats. For instance, firewalls may let through an email irrespective of who sent the email and whether or not it contains a virus.

There are trade-offs between the level of filtering and the need to allow application services to pass through. An attacker may still be able to compromise your internal systems using application traffic, which you allow to pass through the firewall. For example, some worm programs can infect your computer via HTTP, which is the communications protocol used for Internet browsing and web-based applications.

Other security measures

Even with a firewall in place, you still need to take other security measures to protect your internal computer systems, including:

- arrangements to control people's access to the computer system and the types of information they access.
- a means for uniquely identifying and authenticating each authorised user of the computer system, such as a user ID and password, a smart card and PIN, or biometrics.
- audit and monitoring tools to detect intrusion and other forms of misuse.
- regular off-site backup of the system data for disaster recovery.
- physical security to prevent after-hours access to facilities that house the computer system and associated data storage media (CD, disks, etc).
- arrangements for the proper erasure of patient data prior to disposing of obsolete computer hardware and data storage media (i.e. complete data removal procedures).
- virus scanning and SPAM filtering of incoming email.
- encryption services, such as a Virtual Private Network (VPN), to protect communications with other health systems and to allow GPs to securely access their practice systems from home, when visiting patients, or when working from other health facilities.

For more information on practice computer security refer to the *GPCG Computer Security Guidelines and Checklist.*

What if your ISP already provides a firewall?

Find out what your ISP offers

The quality and effectiveness of Internet Service Provider (ISP) firewall services varies from provider to provider. Some ISPs offer firewall capabilities at their network gateway, which sits between the Internet and your access service (referred to as a network firewall). Others offer firewall capabilities in the network access device (router or modem) that is connected your LAN (referred to as a LAN firewall). In the case of network firewalls, some ISPs may provide the same service, or 'single ruleset', for all customers. While this may provide adequate network security it may also restrict your business if the ISP is inflexible about modifying its rulebase to allow you to use non-standard application services. Other ISPs may have a very 'open rule' policy to accommodate all customer requirements, resulting in a lower network security regime, which may not be adequate for your business.

If you are uncertain about the capabilities of your ISP's firewall service, consider switching your service to a provider that specialises in secure services. Some ISPs offer a Defence Signals Directorate (DSD) approved firewall service that meets government standards, albeit at a higher cost.

For more information on Australian Government firewall standards refer to www.dsd.gov.au/infosec/index.html.

Have several layers of firewalls

In line with the 'defence in depth' security principle, it is good practice to have multiple layers of firewall capability at different points in your network. For example, you may wish to use:

- the network firewall service at your ISP gateway.
- a local LAN firewall.
- firewall software on each computer connected to your LAN (personal firewalls).

Personal firewalls should always be considered for computers that are used away from the practice (e.g. laptops), particularly if those computers have remote access into your LAN via wireless or dialup Internet services.

Many ISPs offer fully managed multi-tier firewall services. However, while you may rely on your ISP to provide a network firewall service you may choose to provide your own LAN and personal firewalls. If you are planning to rely solely on your own firewall/s, you need to be confident that you have chosen the right firewall product and that you know how to properly configure and manage the firewall. You should also be aware that some applications, such as video conferencing and Voice Over Internet Protocol (VOIP), may not work ‘out of the box’ through your firewall.

Be prepared to continually manage your firewall configuration, making changes to the filtering rules to accommodate new applications as necessary, upgrading software and firmware with vendor patches to address any vulnerability in the firewall, and monitoring log files for signs of attacks.

Firewall implementation issues

Security versus functionality

Implementation and secure configuration of a firewall may impact on the delivery of some applications services. While a firewall will not limit your ability to allow any services into your network, some services may have less than secure protocols. By allowing these to pass through the firewall you may greatly diminish its effectiveness. Typically, applications that deliver file transfers such as Pathology may use less that secure protocols like FTP (File Transfer Protocol). Configuring your firewall to enable such applications may open up your network to attack through these protocol services.

If your secure firewall configuration is not enabling applications to pass through, you can seek support from the application supplier and a network security expert. Firewall configuration rules may be modified to allow the required application protocols to be accepted, or accepted from an authorised address, or in a particular network direction, so as not to compromise your network.

Alternatively, the application may require an insecure configuration of your firewall. The best advice in this case is to either work with the application vendor to modify the application to provide a secure firewall friendly transfer mechanism, or change to an application that does not require insecure configuration.

Do you have the necessary IT skills in-house?

Recently, network security institution, SANS added a ‘bonus’ number 11 to its list of the worst mistakes by IT people that lead to security breaches: ‘Allowing untrained, uncertified people to take responsibility for securing important systems’.

It is critical that you involve someone with adequate security experience when purchasing and setting up security for your practice computer system/s.

Setting up a firewall always requires basic (inter)networking knowledge. While this tutorial can help you understand more about firewalls and take you through the steps involved, you will still require the appropriate networking and IT security expertise to properly protect your computer system.

If your practice has no skilled IT person on the staff, this tutorial may still act as a guideline for hiring professional IT staff and/or services, and as a checklist for specifying the work that needs to be done.

STEP 1: Understanding firewalls in principle

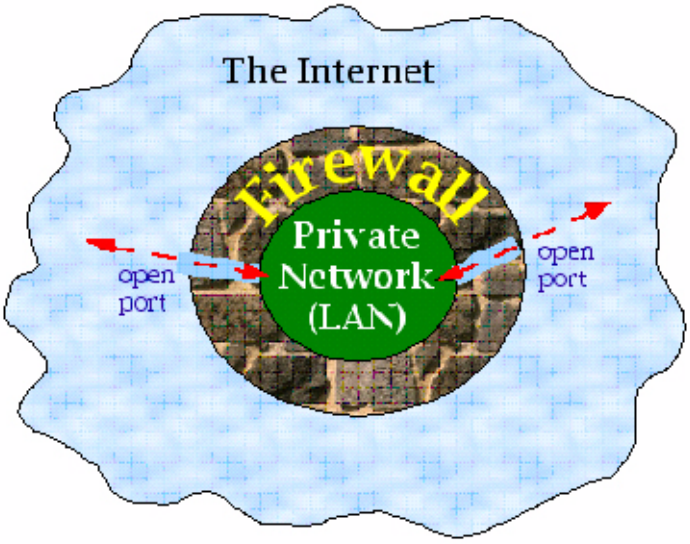

- A firewall is a means of shielding your private computer system from an untrusted network, like the Internet.
- Any outside connection puts your network at some risk, and should be regarded as gateways to an untrusted network, whether or not it is in use. Some standard computer services increase this risk by running less than secure IP protocols such as FTP (File Transfer Protocol) and UDP (User Datagram Protocol).
- Firewalls mediate network traffic to allow authorised traffic and bar unauthorised or risky traffic.
- However, your firewall cannot completely shield you from the outside, as you may want to browse the Internet, send and receive emails etc. For that purpose, your firewall needs to open some doors for traffic between the networks. Step 7, the section on understanding ports and configuration explains which doors are safe to open and how to safely open and close such doors.
- There is no point in establishing a connection to another network if you do not use it.
- If your private network is not connected to any other network, you do not need a firewall.

Understanding that firewalls have limits

- Firewalls are not the panacea of computer security.
- Firewalls are one important tool to secure your network, but are not going to solve all your security problems. They are an essential ingredient in your total IT security strategy, but cannot be the only one.
- In particular, firewalls offer only limited protection from attacks originating from within your private network. They will not prevent you from opening dangerous attachments to emails, or downloading unsafe applications.

STEP 2: Understanding how the Internet works

Understanding how and why to install a firewall, a basic networking knowledge is required. Here is a simplified explanation of how the Internet works, using analogies with the phone system. You need to understand these basics to be able to manage your firewall.

How data is exchanged—TCP/IP and other basics

Nowadays, most networks use a protocol called TCP/IP (Transport Control Protocol/Internet Protocol).

TCP/IP is the collection of communication protocols that manages the exchange of data in what is referred to as ‘packets’. This is the protocol the Internet uses. Whenever the term networking is used in this tutorial, assume that the TCP/IP protocol is used.

To participate in a TCP/IP network, you need a network interface. A computer can contain one or more network interface cards (NIC).

Each NIC must have a unique IP address to participate in a network and to interact with network traffic.

Basic networking principles

Internal versus external networks

Most practices have their own internal phone networks with ‘extensions’. These phones have internal numbers, usually up to three digits, that people within the same practice can use to dial each other.

The internal network works even if the external phone lines are down, because it doesn’t use the external phone lines at all.

Likewise, if somebody from outside the practice wants to call a particular extension, usually they cannot do so directly. Rather they have to ring the receptionist first, and be switched through onto the internal extension. This is done by pressing a special button on the receptionist’s phone, which instructs your practice internal phone switch box to connect the external public phone network temporarily, via the requested phone (extension) in your practice internal network.

Some internal phone systems are installed in a way that allows internal extensions to be dialled directly from the external public phone system. This is convenient for callers, but can also bypass your receptionist, exposing you to calls you might not want to receive.

If you want to call outside the local internal phone network, you have to request a ‘line’ by pressing a special key on your phone so your practice internal phone switch box will establish an outgoing connection for you.

It is similar in the TCP/IP network world:

- telephone company = ISP.
- telephone = network interface.
- private extension phone number = private IP address.
- phone number = public IP address.
- telephone line = network cable.
- PABX/Switchboard = bridge between private and public network.
- receptionist = firewall.

Private IP address

The Internet uses special Internet addressing schemes to distinguish private local networks from computers participating in the Internet.

A computer using any of these reserved addresses, will not be visible directly to the Internet – in the same way your internal phone with its internal extension number cannot be reached directly from the public phone network without your receptionist switching the call through to that extension.

Your ‘private IP address’ is equivalent to the local extension number of your practice internal phone network. However, if an attack breaches your network, all your practice addresses are likely to be exposed to the hacker.

Public IP address

The Internet Service Provider (ISP) connects you to the rest of the Internet via an address that does not belong to these special reserved private numbers. This is your ‘public IP address’. Your public IP address is the equivalent of your official, external phone number – that is, the number which people can dial from anywhere within the public phone network to reach your reception desk.

To be able to use the Internet, a bridge is needed between the private network interfaces and the public Internet. This is the equivalent of your phone switch box (PABX) in conjunction with an arbitrator, your receptionist. The phone switch box makes it technically possible to connect the public to the private phone network, and the receptionist makes sure that no unauthorised caller gets through directly to a specific extension.

In the Internet world, the equivalent of the PABX phone switch box would be a ‘bridge’ or a ‘router’, and the equivalent of your receptionist would be your firewall.

STEP 3: Deciding which firewall product you need

The choice of firewall depends on your needs, based on:

- the risks to your practice information.
- the available IT skills.
- your budget (with some solutions).

Step 8 of this tutorial includes some suggested firewall products that have been reviewed by General Practice testers.

Different firewall scenarios

Depending on how your practice network is set up, and whether or not you want to provide web services, there are different ways of positioning your firewall between the Internet and the computers you want to protect.

1. Simple scenario

The simplest scenario is to place your firewall between your Internet access point (e.g. ADSL modem, satellite modem, cable modem) and your network, as illustrated in Figure 1.

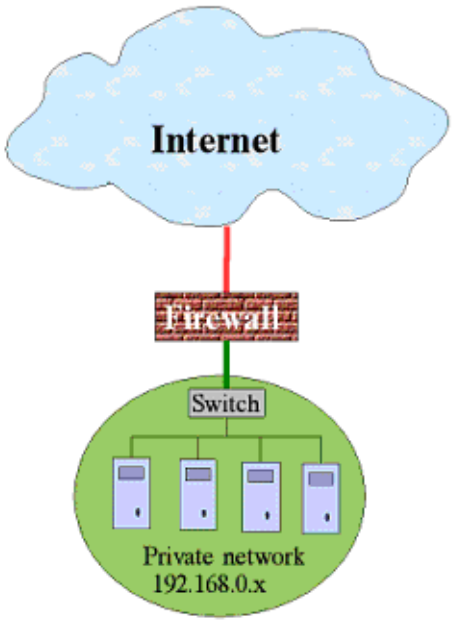

Figure 1: simple separation of public and private network with a firewall

2. Separate web server firewall

If you want to run your own website to advertise your practice and possibly provide information for downloading and browsing, you have to expose the computer that stores that information to the Internet.

That computer does not need to communicate with the rest of your private network, so you can isolate it completely. However, you cannot leave it completely without protection. You still want to protect it from being hacked into and defaced, so it needs a firewall too. Figure 2 is illustrates how to configure firewalls in such a scenario.

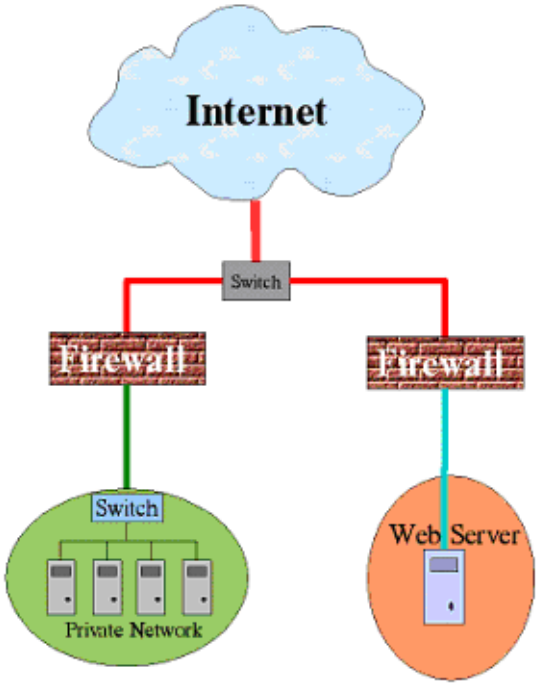

Figure 2: protecting your private network, and protecting your web server with two separate firewalls in two independent networks using a single Internet connection

3. Built-in web server firewall

The poor man's solution to the previous scenario, which is still viable in most circumstances, is to implement the second firewall directly on your web server. However, it is most likely that it will be more expensive and harder to administer than simply putting a $100 box in front of your web server, as a 'built-in' firewall.

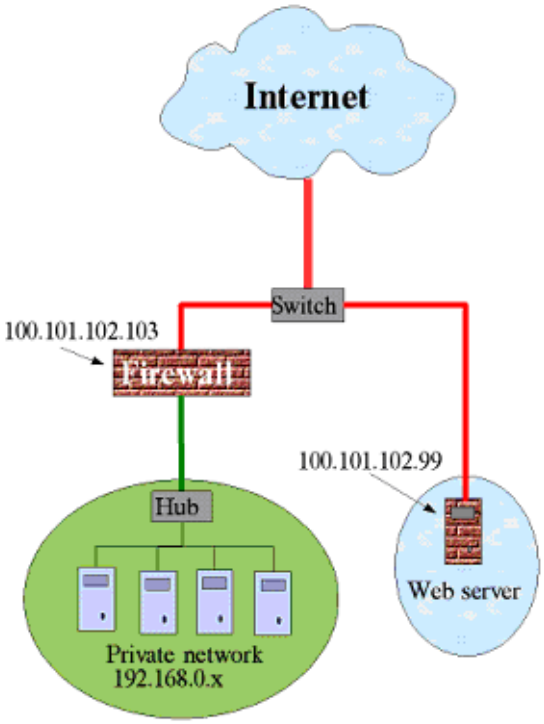

Figure 3: protecting your private network with a firewall, and protecting your web server with a 'built-in' firewall in two independent networks using a single Internet connection

4. Web server as separate (perimeter) network

Once you want to provide web services to the outside world, such as online appointment bookings, you will probably need a slightly different layout.

Computers that are exposed to the outside (the Internet) for access – and that includes remote access for maintenance purposes etc – should be placed into a separate ‘perimeter network’ sometimes called a Demilitarized Zone (DMZ).

Computers placed in the perimeter network can usually communicate with computers within the private network but only in very strictly controlled ways.

Some firewall devices provide separate network interfaces for this purpose and they can manage the private and perimeter network in different, adequate ways.

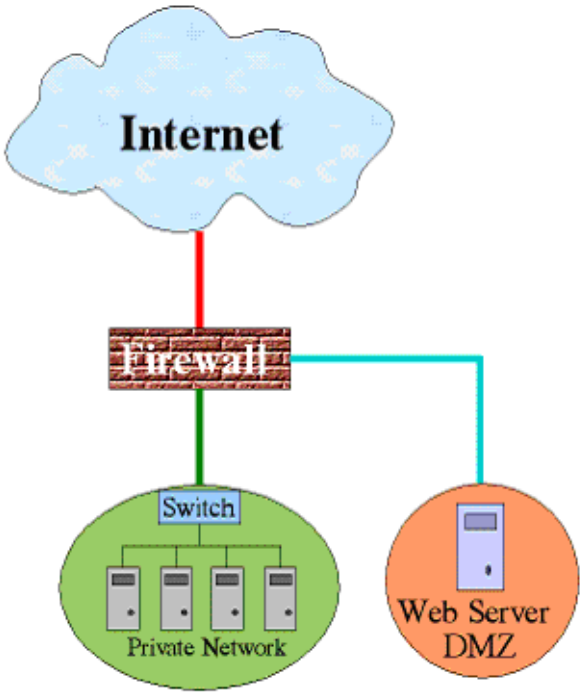

Figure 4: A single firewall handling both the private and the perimeter network

STEP 4: Understanding firewall technologies

Simple versus sophisticated

Once you have decided on the general network layout and where to place the firewall, you have to think about what firewall technology to use. Unfortunately, there is no simple right answer that covers every circumstance.

Firewalls can use simple or sophisticated methods to do their job. More sophisticated firewalls are usually safer if properly configured, but configuration can be much more difficult.

Rule number one – a properly configured simple firewall is more secure than a poorly configured more sophisticated firewall.

This is important to understand. Do not aim for highly sophisticated devices if you do not have the expertise (or an expert) to set them up and maintain them.

Even simple packet filtering firewalls can achieve sufficiently secure separation of private and public networks in a General Practice environment, as long as they are properly configured. Similarly, the way a firewall is implemented in your local network is critical.

Rule number two – the firewall should be the only entry/exit point in your network. If not, you potentially have an open backdoor in your network.

Here is a simple analogy. Imagine you have to defend a narrow passage into a castle. You can choose a simple heavy club as a weapon, or a sophisticated pistol. While the pistol at first seems the better choice, you might discover that the club will never fail you and still do the job in most cases, while you cannot really predict when the pistol will fail, you need training before you can use it, as well as ammunition. If you are not experienced with pistols, you are probably better off with the simple heavy club.

NAT (Network Address Translation)

This is not really firewall technology; rather it is a prerequisite for separating private and public networks. However, you should be aware that some standalone NAT products on Windows (e.g. Windows 98 Internet Sharing) are advertised as firewalls even though they are not.

NAT is essentially a mechanism to route traffic from a private network addressing scheme to the public Internet addressing scheme. Anything behind a NAT router is already difficult to reach from the outside.

How NAT works

Imagine a practice with an internal phone system of four phones. They have the internal numbers 1, 2, 3 and 4. If somebody dials ‘2’ from any internal phone, they will be connected to extension number ‘2’. But anybody outside, from the public phone system, dialling ‘2’ will not be connected to that phone. Why not? Because the public phone system uses a specific phone number system which is different from the internal phone number system.

However, anybody can dial the public phone number of that practice and the receptionist can put the caller through to extensions 1 to 4 if requested and if it is appropriate.

The NAT router does essentially the same job as the receptionist, translating your own internal network addresses into public networking addresses and vice versa.

Sideline: Early 2004, my home firewall died suddenly. It was temporarily replaced with a simple NAT router ('e-smith Linux distribution') to distribute the single dial-up Internet account to the whole family. An old version was installed since we could not access the Internet to download the latest one. I was called away to a patient just at the end of the installation and when I came back I was too tired to download and install all the security patches that were made available since that version was released. By the next morning, that NAT router had already been hacked. Fortunately, nothing but a sacrificial honey-pot computer was connected and no confidential data was accessed or lost.

Lesson learned: not even dial-up lines are safe, and security patches definitely cannot wait overnight. Never go live with an untested and un-configured firewall.

Dr Horst Herb

STEP 5: Understanding different types of firewalls

Packet filters

Data transferred via TCP/IP protocol is usually sent in the form of ‘packets’.

Each packet contains small amounts of data attached to a ‘header’ which has information about the purpose, source and destination of the packet.

A packet filtering firewall looks at each packet and, depending on a nominated set of rules, decides whether or not to let them pass through. It filters packets depending on rules set about the port, direction (ingoing/ outgoing traffic) as well as source and destination IP addresses. Using the analogy to the practice phone system – a doctor’s receptionist receives many calls but doesn’t automatically put them all through to the doctor. The receptionist only puts through calls that are agreed as appropriate to forward to the doctor. All other calls are blocked. (Alternatively the receptionist can forward all calls except those on a ‘denied’ list).

Advantages of simple packet filters

- Very fast – no bottle neck caused by the firewall.
- Use few resources – devices are inexpensive, draw little power, and generate little heat.
- Simplicity – less chance for faulty implementation (bugs), often more robust than more complex solutions.

Limitations of simple packet filters

- Formally invalid TCP packets or packets that do not seem to belong to an active connection cannot be filtered.
- UDP packets cannot be filtered properly. This is an important drawback. You can either block all UDP transactions or accept that you are vulnerable in that regard.

Working around the limitations of packet filtering

- Always make sure that you apply security patches as soon as they become available to all network active applications. Attacks that slip though a packet filter based on incorrect packages will only work as long as there are any exploitable faults in application networking code.
- Block all incoming UDP connection attempts. Most users in General Practice use UDP only for DNS lookups and will not provide DNS services themselves.

Examples of packet filtering firewalls

- Older Linux systems (Kernels 2.2x and lower).
- Many routers (e.g. CISCO ACLs, most ADSL modem routers / wireless routers).

Stateful packet filters

This is essentially a packet filter that knows about the history of a packet and can see it in the context of a connection.

Using the analogy to the practice phone system – somebody rings and asks to be put through to Dr X. He claims he is returning a call. A ‘stateful’ receptionist keeps a phone log and checks if Dr X did ring that person in the first place, and will only put this call through if the log indicates that this is a legitimate return call.

Advantages of stateful packet filters

- Simple and fast technology.
- Protects against ‘answer’ session exploits.
- Protects against some DoS attacks like ‘SYN flooding’.

Disadvantages of stateful packet filters

- Vulnerable to attacks with malformed packets (since it does not know about packet content).
- Vulnerable to protocol-based attacks / ‘buffer overflow’ attacks.

Examples of stateful packet filters

- Linux NetFilter based firewalls.
- BSD IPF or OpenBSD IPF based firewalls.
- Watchguard Firebox.

Stateful inspection packet filters

This is a stateful packet filter armed with protocol specific modules that actually know how to interpret a packet in the context of its protocol. Also known as dynamic packet filtering, stateful inspection provides enhanced security by keeping track of communications packets over a period of time. Both incoming and outgoing packets are examined. Outgoing packets that request specific types of incoming packets are tracked; only those incoming packets constituting a proper response are allowed through the firewall.

In the practice phone system analogy – as with the ‘stateful’ receptionist, the ‘inspecting stateful’ receptionist will only accept calls from patients who are confirmed to be returning a call. Imagine though, that the inspecting stateful receptionist puts the call through to Dr X but then listens in so that if the patient starts to ask about unrelated problems. If this should happen, the receptionist interrupts the connection and explains that the patient will need to make the relevant appointment.

Advantages of stateful inspection packet filters

There are the same advantages as stateful packet filters (above) plus:

- protection against some protocol based attacks.
- less vulnerable to misuse of open ports.

Disadvantages of stateful inspection packet filters

- Depends on protocol specific inspection modules. Protocols not covered by inspection modules will be handled no better than with a stateful packet filter.
- Needs a lot more processor power and RAM, hence is more expensive, generates more heat and is more prone to technical faults.

Examples of stateful inspection packet filters

- Firewalls based on newer Linux kernels (2.6).
- Sonicwall appliances (for limited number of protocols).
- Checkpoint Firewall appliances.

Application proxies

This type of firewall goes one step further than stateful inspection firewalls.

It not only knows the history of the connection, but also inspects the data within the packets, and decides whether or not to allow a packet pass through depending on the content.

The proxy basically receives packets, analyses them, and repackages them safely according to nominated rules before sending them on as instructed.

Application proxies are typically located within a separate ‘perimeter network’ or Demilitarised Zone (DMZ) that is a third network between your real private network and the public Internet.

It insulates the internal network by enabling less secure services to operate in the perimeter without compromising the internal network.

Advantages of proxy type firewalls

- Protection against malformed packets.
- Protection against more protocol based attacks than stateful inspecting packet filters can provide.
- More granular control over which protocols will traverse the networks.

Disadvantages of proxy type firewalls

- Rather complex – needs more powerful hardware and therefore generates more heat and is more prone to technical faults.
- Due to complexity of the software, is more likely to contain programming errors (‘bugs’).
- A specialised proxy is needed for every single protocol – you may need custom written software to proxy some of the networking applications you use.

Examples

- Tinyproxy, Squid, Exim, Sendmail, Smtpfwdd.

Full-blown application proxies exceed the scope of this tutorial. Organisations that are active enough on the Internet to need them should employ professionals who are fully experienced in this field.

Compromise solution

However, there is a simple compromise where you can gain some of the benefits of full blown application proxies in a perimeter network, through a little extra work (plus an extra network interface and one extra dedicated computer):

- A single computer is connected via a separate network interface using a separate address range.
- This computer hosts a small number of applications (like web server and email server) that are allowed to communicate with the public Internet without putting the private network at risk.
- A separate ‘perimeter network’ or Demilitarised Zone (DMZ?) is very useful for practices that want to provide email and web services, accessible from the public Internet, without having to outsource these services.

In the analogy to the practice phone system – a person rings to discuss a test result with Dr X. Instead of disturbing the doctor (and risking discussion about other unrelated issues with that patient), the receptionist puts the caller through to the practice nurse who knows all about the results, but has no details about any other matters. Thus, even if the caller is an impostor, they will not gain any additional knowledge about a patient’s confidential records other than perhaps the test results.

STEP 6: Understanding network addressing

The first question before you set up your firewall will always be: what address range are you using in your local network?

Currently, Internet addresses are unique 32 bit numbers, usually displayed for better memorability as four 8-bit numbers separated by full stops (that is anything from 0.0.0.0 up to 255.255.255.255).

Some of these many possible addresses are reserved for special purposes, like local area private networks.

Local Area Network (LAN) addresses

Local area networks (LANs) are supposed to be private networks separated from the public Internet. To allow the same protocol for the public Internet as for the private network (or intranet), several blocks of possible addresses have been reserved for private use.

The address reserved for private local area networks are:

- 10.0.0.0 to 10.255.255.255.
- 172.16.0.0 to 172.31.255.255.
- 192.168.0.0 to 192.168.255.255.

These addresses will never be visible to the public Internet, unless the addresses are deliberately translated in to public addresses first.

The most common way of address translation is called NAT (Network Address Translation). Simple routing would not work, because nobody from the outside (Internet) can contact any private address without it first being translated into a public address.

Every single address used in the public Internet must be unique (i.e. the same public address is never used for two computers anywhere in the world). However, there are an unlimited number of private subnetworks using the same addresses. But within each private Intranet, the same rules apply as within the public Internet: every computer must have a unique address.

Which address block to use for a medical practice network

To avoid confusion, this tutorial will refer to network address blocks that always start with 192.168 when citing examples.

Subnetworks

The third number in the string of numbers (i.e. the digits after the second full stop) will be specific for your ‘subnetwork’.

You can have multiple subnetworks within your practice or practices, but only machines within the same subnetwork will be able to see each other without a special bridging interface.

For example: Your subnetwork address is 192.168.0.

Now, we can identify up to 256 different network interfaces (and one computer can have multiple interfaces), namely 192.168.0.0 to 192.168.0.255.

Subnet masks

To allow all these 256 interfaces to see each other, you have to specify a subnet mask.

In the example above, if we want all of the possible 256 computers to see each other, you need to specify the subnet mask as: 255.255.255.0.

Imagine 8-bit (0 – 255) as 8 little switches. Each switch that is ‘on’ has to be matched by the corresponding switch in your address. The number 255 hence indicates that an exact match is required. The number 0 represents the other extreme: all possible 256 numbers (0 - 255) would match. Thus, a subnet mask of 255.255.255.0 would allow all IP addresses ranging from 192.168.0.0 to 192.168.0.255 to ‘see’ each other.

If you need to restrict your subnet further, or expand it, you will require more knowledge than we can cover in this brief introduction. You can start further reading at IP Addressing and Subnetting for New Users.

DHCP servers

In larger networks, or in networks where new devices are added regularly (like wireless networks), allocating a correct and unique IP number to every network interface can be a daunting task. Whenever you decide to change your subnetwork, for whatever reason, you also need to change all connected interfaces.

It is easier to leave the task of allocating IP addresses to a DHCP server. In this case, you can set up all your client computers just once to automatically get the networking information from a DHCP server, and then you can forget about them.

The following links provide more information on how to set up a network interface to get its address from a DHCP server:

- if you are using Windows.
- if you are using Apple Macintosh.
- if you are using Linux.

For other operating systems, please consult your manuals or engage an IT professional who knows the system.

To set up the DHCP server, you first have to know where it is:

- if you are using an embedded gateway/firewall (e.g. D-Link, Linksys, Billion, Netgear), go to Embedded DHCP Server Setup.
- if you are using a dedicated gateway/firewall computer (e.g. Smoothwall, SME Server), go to Dedicated DHCP Server Setup.

Setting up a DHCP server usually requires answering the following questions:

- what is the DHCP server’s address?
  - in our example: 192.168.0.1
- what is your start address?
  - in our example: 192.168.0.50
- what is your end address?
  - in our example: 192.168.0.250
- what is your DNS server’s address?
  - find out from your ISP

STEP 7: Understanding ports and firewall configuration

What are ports?

Ports are special addresses within a network address that are required to access various network services. For example: your address is 1.2.3.4, and you want to access the web server – choose port 80. If you want to access your POP3 email server instead, select the same address but choose port 110.

Here is an analogy: You have an office building with a street address, e.g. 8 Smith Lane, Melbourne. With this address, you can find the right building. But this building is 50 storeys high, with many offices at each floor. Some offices may be open for public access, and some may be closed to the public. To find the right office, you have to specify the floor and office number. In many cases, a receptionist will only let you through if you can specify exactly what floor and office you want to visit.

It is similar in networking. A TCP/IP network address allows you to identify a specific network interface (usually a computer), but that is all. To access a specific service (like web browsing, sending and receiving email, transferring a file with FTP) you have to specify the port you want to contact. Most computers are configured to have their web server reply as ‘default port’ whenever the computer is contacted without specifying what service is wanted.

Which one to keep open, which one to close

When you configure a new firewall, initially close all ports. In most cases, there will be no reason to open up any.

The first step should always be:

- close all ports using the firewall ruleset configuration functions – this will require knowledge of the firewall ruleset syntax requirements.
- check whether you can access all Internet services you need from your connected local area network computers. If everything works, don't open up anything.
- now run a Firewalls: Audit scan over your network, e.g. ‘Shields up’, and see whether your firewall holds.
- ensure there are no other un-firewalled network connections.

Obviously, you cannot totally shield your own private network from the Internet if you want to access Internet services, like email and web browsing. So you will selectively open some ports in your firewall if needed. However, in most cases if you don’t provide a service to the Internet yourself (e.g. a web or email server), you can keep the ports closed because your firewall will allow connections that you have initiated from your side.

Well known ports

The 131 070 (1 - 65535 possible ports for both TCP and UDP) are divided into 0 - 1024: ‘well known ports’. These are usually public Internet services like HTTP, FTP, SMTP, POP3, IMAP etc. Of these, 0 - 255 are the ‘registered TCP/IP services’ which are platform independent, whereas 256 - 1023 usually represent UNIX services (Windows and Macs weren’t around when the Internet took off, and both took a long time to adjust to the concept TCP/IP networking).

Reserved ports

‘Reserved ports’ are 1024 - 49151. These are usually registered with the Internet agency IANA, to avoid conflicts between software products.

Public ports

‘Public ports’ are 49152 - 65535. These are up for grabs so never rely on these ports delivering the same service.

However, this is all entirely voluntary. Nothing stops you from running your web server using port 21 instead of port 80, although it would not be sensible to do that. It is worthwhile remembering that writers of malicious code (backdoors, Trojans etc) do not have to follow convention regarding the port numbers they use. Nor will it help you to try to disguise a web service behind an unusual port number. Most good port scanners not only detect open ports, they also find out what protocol is available through that port.

IANA has a comprehensive port list. Of those, the ones you are most likely to need are listed below. Remember that through these ports a service is requested, that is they have to respond on the server side. Most protocols allow the requesting client to choose an arbitrary port above 1023 to communicate with the server. Hence, usually you don’t need to open these ports on the client side.

Workstations (computers that are not servers) can usually close all ports for incoming traffic without compromising functionality. This means you can keep all incoming ports closed as long as you do not operate a server (for external web services) behind the firewall.

Ports you are most likely to use

- Port 22 (SSH) – secure shell access, for remote management purposes, remote backups etc. Keep closed if you do not need it.
- Port 53 (DNS) – name services, translating address like http://www.gpcg.org.au into a valid TCP/IP address; outgoing only, close for incoming packets.
- Port 80 (HTTP) – world wide web. Open it if you operate a web server, otherwise close for incoming packets.
- Port 25 (SMTP) – transfer of emails.
- Port 110 (POP3) – email server using POP3 protocol.
- Port 143 (IMAP4) – email server using alternative IMAP protocol.
- Port 995 (secure POP3) – instead of port 110, if your mail server supports the POP3 protocol over TLS/SSL.
- Port 993 (secure IMAP4) – instead of port 143 if your mail server supports the IMAP4 protocol over TLS/SSL.
- Port 443 (SSL) – secure socket layer; web browsing or email exchange via secure encrypted link.

Unless you are an expert who knows exactly what to do, you should not have any ports open in your firewall other than the ones listed above, and you should only open those if they really need to be opened.

Once again: Close all ports, try your applications out, and only then, if something does not work, try opening that port (for outbound traffic only first, if your firewall configuration options allow this).

STEP 8: Suggested firewall products

Your choice of firewall depends on both your needs and IT skills, and with some solutions it also depends on your budget.

Government Security Adviser (DSD) recommendations

The Australian Government also provides recommendations and advice on firewall products, through the Governments Security Advisor, Defence Signals Directorate (DSD). DSD provides a list of evaluated products that are certified under a rigorous evaluation program, where the security claims of the vendor are tested using international evaluation criteria. These products are applicable to larger organisations and as such meet higher-end security needs and cost more. The most up-to-date DSD approved product listing is at www.dsd.gov.au/infosec/evaluation_services/epl/dap.html

There is also useful tabular overview of firewalls / routers commonly available in Australia available at www.ozcableguy.com/quickref.html.

General practice tested solutions

The solutions suggested in this tutorial are limited to products that have been tested by General Practice experts, who have provided 'peer reviews' of the products and solutions. There are many other viable solutions available on the market, but the GPCG cannot comment on suitability nor give advice on configuration of these. This tutorial only suggests products where commercial support is available. (This does not mean that non-commercial products are inferior, but using them requires such a high level of expertise that people able to use them without risk would not require the help of this tutorial.)

The general recommendation for any practice would be a router/firewall appliance with:

- Failover / load balancing dual WAN (Internet) ports – meaning it can connect simultaneously to two different broadband providers (e.g. ADSL + Satellite, ADSL + different ADSL provider, ADSL + Wireless, cable + ADSL). That way you can not only improve the performance of your Internet connection substantially but if one connection fails, the device will automatically re-route all traffic to the other provider. If you depend on the Internet – and most of us will depend on it sooner or later – this is one indispensable feature.
- Good logging facilities – no firewall is perfect, and you have to know when you are under attack. Most people will rarely inspect cryptic log files that are a hassle to access, but some firewall appliances will email you such logs in understandable form
- VPN capabilities – sooner or later you will discover the convenience of being able to access your surgery remotely via the Internet. Of course, such connections have to be as secure as possible. Quality products rely on the IPsec standard for such virtual private tunnels through the Internet. They automatically establish a secure link between two connecting appliances, and allow you to use any IPsec compliant client software on your own computer when you cannot connect from one device to another directly (e.g. between branch surgeries with two compatible devices installed).

Unfortunately, not many devices meet these recommendations. Here are a few suggestions according to price:

1. under A$500 – Netcomm NB740
2. under A$600 – Linksys RV082
3. under A$1 200 – Zyxel ZY70
4. under A$2 000 – SonicWall TZ170E

There is a product for less than A$200 which appears to fulfil all criteria, but it doesn't have an Australian distributors yet (and is yet to be tested by General Practice testers). It is the Hawking H2WR54G.

Products costing more than the A$600 solution offer some additional features. Always study the product information thoroughly before making a decision.

The list above is not exhaustive – it represents products reviewed up to now.

After you make your choice and install your firewall, refer to this tutorial's checklist before you connect your private network to the firewall.

If none of the solutions suggested here suit you, there are further firewall options below, listed according to security requirements.

Low security need/low IT expertise

This type of medical practice would use a single means of connecting to the Internet (e.g. ADSL or dial-up). Computers storing confidential information would not be connected directly or indirectly to the Internet. For example, a single, non-networked computer is used for email, web browsing and all other Internet transactions

Recommendation: embedded gateway device (modem plus firewall in one small dedicated box).

Reviewed products (in alphabetic order):

1. Billion products
2. D-Link products
3. Draytek products
4. Dynalink products
5. INEXQ products
6. Linksys products
7. Netcomm products
8. Netgear products
9. SMC products
10. Snapgear products

Why not just install firewall software on the Windows box (like ZoneAlarm or Norton Internet Security)?

A chain is only as strong as its weakest link. It is the same situation with firewalls – they can only be as secure as the underlying operating system and the software running on that system. While there might not be anything wrong with such firewall products in themselves, other programs running on your PC may bypass or negate the firewall function. There is no point running a firewall application on a computer that cannot itself be secured.

Considering that some of the router appliances listed here cost less than $A70 (street price), and already include stateful packet inspection technology, it would be advisable to have at least one of them. Of course, you can use your preferred firewall software on your Windows box as well if you wish.

Intermediate security needs/intermediate IT expertise

This type of medical practice would typically grant Internet access to all desktop computers. Doctors would be able to browse the web and use email from the same computer they use for health records. Most practices would fall into this category.

Recommendation: dedicated gateway computer with firewall and automated security updates or a firewall appliance with customisation/upgrade features. If budget and expertise allow, you can also chose from the 'high security need' category (recommended).

Reviewed products with commercial support:

1. Gibraltar
2. SME Server and Gateway
3. SmoothWall

Reviewed embedded devices/firewall appliances:

4 Billion products
5 D-Link products
6 Draytek products
7 Dynalink products
8 INEXQ products
9 Linksys products
10 Netcomm products
11 Netgear products
12 SMC products
13 Snapgear products

Reviewed products without commercial support:

14 DevilLinux
15 Euronode
16 IPCop
17 NetBSD Firewall Project
18 RedWall
19 Sentry Firewall CD

High security needs/high IT expertise

This practice does not only grant Internet access to all networked computers, but it does also provide web services (e.g. practice website, online appointment system etc) from within the practice premises, and/or it uses the Internet for virtual private networking VPN – e.g. a link between branch surgeries). This practice will probably operate its own mail server.

However, even practices with lesser security demands would benefit from firewalls of this category. Both software products and appliances in this category often have additional features, such as more than one WAN (Internet) port, allowing you to balance multiple broadband connections for better performance and failover capabilities.

Recommendation: Dedicated customised gateway computer with ‘DMZ’ and VPN tunnelling capabilities and stateful packet inspection firewall. It is strongly recommended that you choose products with good commercial support unless you have outstanding IT security expertise within your staff. Alternatively, more sophisticated (and expensive) firewall appliances may be suitable. Choosing a product with load balancing/ failover capabilities over at least two WAN (Internet) ports is highly desirable if you depend on availability of Internet services and can afford multiple ISP services.

Recommended products:

1 Astaro
2 CISCO PIX Series
3 Cyberguard SG570 – this model no longer in production
4 Linksys RV082
5 Securepoint
6 SonicWall TZ170E
7 Watchguard Firebox

STEP 9: Principles of firewall configuration

This section explains the steps necessary for configuring any firewall. Product-specific information is in the section on ‘suggested firewall products’.

To configure your firewall, you may need to connect it to a computer.

- Make sure that the configuring computer is not connected to any other computer (e.g. via wireless connection) – only one network connection is allowed, and this is between the configuration computer and the firewall.
- Read the firewall manual regarding the default IP address your firewall will have. Configure your configuration computer so that it can establish a TCP/IP connection to the firewall – the manual usually tells you how to do this.
- Follow the instructions in your manual to open a web browser in your configuration computer, and point it at the firewall's configuration address. Quite often that will be http://192.168.0.1:80 or http://192.168.0.1:8080. Usually, you will be prompted for a user name and password – consult your manual for the default password.

The first thing to do after logging into your firewall is change the password.

- It is important that you chose this password well – anybody who could access your firewall configuration interface from the Internet could take over your whole network if your password is too easy. Read how to choose a good password before you chose one.

After you have changed your password and logged into the firewall configuration interface using your new password:

- Close all ports on your firewall using the firewall’s configuration interface. Do it before you do anything else. If you cannot find it in the interface menu, consult your manual. It only takes a couple of mouse clicks to do so. Some devices have all ports disabled by default, and this is how it should be.
- Now configure the firewall’s WAN (Internet) interface via the configuration screen. In most cases, the ISP will have provided you with a static IP number and DNS server details which you will enter. Alternatively, some have their IP number allocated dynamically. In that case, set the firewall’s WAN address to ‘DHCP client’ or ‘automatically allocated’. Many firewall devices have to restart after this.
- Connect your firewall device or system to your Internet connection (ADSL modem, cable, satellite modem, ethernet port). At this stage no other computer connected to your Intranet is allowed to be connected to your firewall. Still, only the configuration computer is connected to your firewall, plus the ADSL or cable or satellite modem plugged into Internet (WAN) port of the firewall.
- Now open another browser window on your configuration computer. Try to browse the web, e.g. http://www.google.com. If it works, try to send and receive a test email. If that works, you do not have to open any additional ports on your firewall.
- The last step is to test your firewall. You need a few more skills and tools for that and some hints are given in the section Step Ten – DIY security audit Take this task seriously. If you cannot do it yourself, you must contract somebody who can before you consider connecting your practice network to the Internet. If you do not test your firewall, your patient's confidential data stored on your network is at risk.
- If you are happy with your setup, run through the tutorial’s final checklist.
- Now you can connect your practice network to the firewall!

STEP 10: DIY security audit

How to find out if your firewall really works

The proof of the pudding is in the eating. The proof of your firewall is in withstanding attacks. Performing a thorough security audit is a task best left to qualified professionals. However, there are some tests that anybody can perform. This section will help you chose security audit products, and understand how to use them and how to interpret the results.

Online testing facilities

The simplest way to test your firewall is to use your web browser from within your private network (behind your firewall) and log in into some of the browser-based security auditing services.

Free test through Shields up!

Steve Gibson from Gibson Research Corporation (GRC) provides his testing suite Shields Up! free online. Just follow the links to Shields Up! on the GRC website.

NMap scan via the web

NMap is a useful scanning tool. Some of its functionality is available at www.whatsdown.net/nmap.html.

Free scan through AuditMyPC

AuditMyPC reports are simpler than the Shields Up! facility, so if those reports were confusing, try these for a start. However, if they don't show up any problem, do **not** assume there are none. This test is very superficial only. AuditMyPC is at http://auditmypc.com/freescan/scanoptions.asp.

Others – be aware!

A variety of commercial providers advertise free online security scans, like Symantec or Sygate. However, the many trials run for this tutorial did not extract any meaningful results from them.

Zonealarm's online offer prompts you to download and install an ActiveX control in order to perform the scans. This is not recommended as any meaningful security policy will invariably prohibit the download and installation of ActiveX controls.

Security audit software

Running auditing software really only makes sense if you can trust the platform you use to run the software. This is why most of the best auditing tools typically run on Unix based operating systems (like Linux or BSD Unix), but some run on Mac OS/X as well. Very few useful and trustworthy security auditing tools run on Windows – and most come at a very steep price.

Local Area Security Linux

This is a valuable tool chest of network auditing and forensics applications that can be run from CD without needing to install anything. You simply boot your computer from this CD. You can also install it on a USB key instead of a CD, and boot from the USB key if your BIOS supports it. Those familiar with Linux will not be surprised that this 200MB system also doubles up as emergency router and stateful packet inspecting firewall. (All it needs is a computer able to boot from either CD or USB port, two network cards, and a few minutes to boot and configure.)

Phlak

Another self-booting CD with a huge number of security related tools, including those needed to test your firewall. If you already have basic networking knowledge, try it out. Although it is based on Linux, it comes with a variety of MS Windows related applications and tools. However, most of the networking security related applications that come with PHLAK are independent of the system and platforms you use.

Most of the tools below will run only on Unix like computers. However, if there are only Windows computers in your practice, a reasonably small download will provide you with everything you need to install a quite sophisticated auditing system on any spare computer within minutes.

Sentinix

Before you download, read the step-by-step installation guide to make sure you will be able to do it. Alternatively, you might want to try it out first before installing anything. You can trial it at http://sentinix.org/demo.shtml. Please read the instructions carefully before you click and write down all the user names/passwords you'll need to try everything out.

SATAN **(Security Administrator Tool for Analysing Networks)**

Satan is the grandfather of most network security auditing tools. It was written in 1995, and is now a bit outdated compared to the newer tools like NESSUS, SARA or SAINT. However, it is free, it still works, and its source code remains a solid baseline for more modern tools of this class.

NESSUS

Nessus has become the Swiss Army knife of network security auditing in experienced circles. It is one step up from all the other tools mentioned in this section. Unlike many other security scanners, Nessus does not take anything for granted. For example, it will not simply assume that a given service is running on a fixed port. If you run your web server on port 1234 instead of the standard port 80, Nessus will still detect it and test its security. It can also be scripted to run customised attacks easily through its built in Nessus Attack Scripting Language (NASL).

Unfortunately, as with most professional tools, it will not work right out of the box. To make best use of Nessus, it requires configuration. Fortunately, it comes with excellent documentation and step by step instructions.

SARA **(Security Administrator's Research Assistant)**

A good starting point to this comprehensive and reliable tool is the training slides that are available on the web at www-arc.com/sara/overview/sld001.htm. SARA (and NESSUS) is already installed and configured ready to run on the bootable Local Area Security Linux CD – there is nothing to install, just boot from CD.

Nmap

Nmap (Network Mapper) is a free open source utility for network exploration or security auditing. It was designed to rapidly scan large networks, although it works fine against single hosts. Nmap uses raw IP packets in novel ways to determine what hosts are available on the network, what services (application name and version) those hosts are offering, what operating systems (and OS versions) they are running, what type of packet filters/firewalls are in use and dozens of other characteristics. Nmap runs on most types of computers and both console and graphical versions are available. Nmap is free software, available with full source code under the terms of the GNU GPL.

SAINT

SAINT is one of the top ten SANS certified security auditing tools. It is not to be confused with the free network monitoring tool: NetSaint. Free trial versions of SAINT are available.

Other useful tools for network security

Honeyd

This is a useful tool that allows you to create complete virtual networks on a single computer. A bit like a reverse Trojan horse, you create this appealing and less secure network behind your firewall, hoping that the intruder will fall for this honey-pot rather than digging deeper and discovering your (much harder to get into) real local network. It can be most educational to run such a honey-pot without a firewall – sometimes it takes less than an hour before the first intruders find their way in and start wreaking (fortunately, only virtual) havoc.

Suggested reading: Honeypots - Sticking It to Hackers, published in the reputable Network Magazine, or the more specific article Bait and Switch with Honeyd which requires more background knowledge to understand.

STEP 11: How to audit your firewall – step by step

Prepare your test scenario

You need two computers plus your pre-configured firewall. We will call the attacking computer (from the simulated Internet) **Nessus**, and the victim (our private network) **Honeypot**.

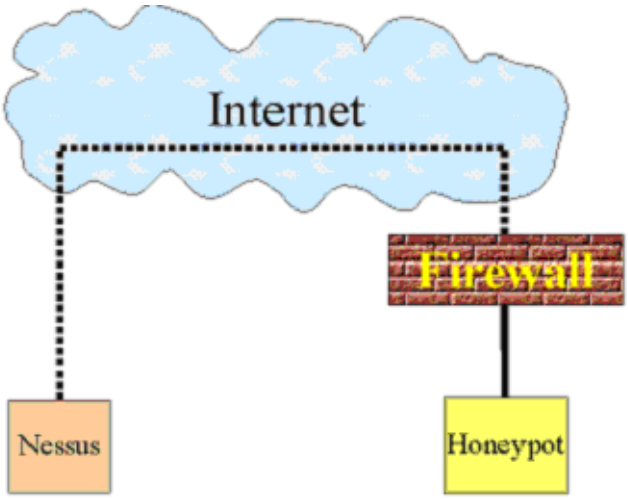

Figure 5: testbed for firewall auditing

Both Nessus and Honeypot will be exposed to the Internet without firewall protection (we have to assume the worst); hence you must not use any computer for this purpose that contains any confidential data.

1. Nessus needs to be connected to the Internet:
   - You can use the same broadband connection as the one your router is connected to, if you plug Nessus into a hub between your ADSL modem (or whatever Internet connection you have) and your firewall.
   - If you have more advanced IT skills, you can simulate the Internet for testing purposes.
2. Recommended configuration for Honeypot:
   - Operating system – whatever you typically use in your practice.
   - Software – install all software you are using in your practice on any computer connected directly or indirectly to the Internet (e.g. MDW, Pracsoft, all pathology download applications, email client). Of course, you will not install any live patient data – most software comes with demo data sets, so use these. Reminder: no software should be installed on any of your private network computers prior to testing on Honeypot.
   - Test configuration – make sure you can browse the web and send and receive emails from Honeypot if this is what you usually would allow in your normal practice setting.
3. Do not install any software on the Nessus attacking computer. Boot the computer from CD without touching the hard disk of Nessus at all – in fact, it will work just as well if there is not even a hard disk attached to Nessus.
4. Recommended configuration for Nessus:
   - Any computer that will boot from CD will do.
   - Configure the computer BIOS to boot from CD, usually by pressing the F2 or Del key while booting, and then entering a BIOS configuration menu (if you don't know how to do this, you will need professional help to do a firewall audit).

- Download PHLAK – this is a 400+ Mb large ISO CD image. Use your CD burning software to create a bootable CD from this ISO image. It **will not** work if you just copy the ISO file onto a data CD.
- Boot from your PHLAK CD while connected to the Internet (if you have dial-up Internet only, you can connect after booting PHLAK).
- If you booted PHLAK successfully on Nessus, you will see a menu bar like this one:

- Select the left hand terminal window icon (on the screenshot above circled in red).
- In the terminal that opens, type: ‘nessus-mkcert’ (without the quotation marks) and answer the onscreen questions.
- If you have generated the certificate successfully, now type ‘nessus-adduser’ (again without the quotation marks). Make up a user name and a password in order to answer the onscreen questions.
- If you have successfully generated a nessus user, you can start the Nessus daemon: type ‘nessusd-D’ (without the quotation marks).

Run the test

1. Open another terminal window by again clicking the left hand terminal symbol on the menu bar at the bottom of the screen.
2. Type ‘nessus’ (without the quotation marks).
3. A graphical user interface will pop up. Enter your user name and password (as used during the ’nessus-adduser’ process).

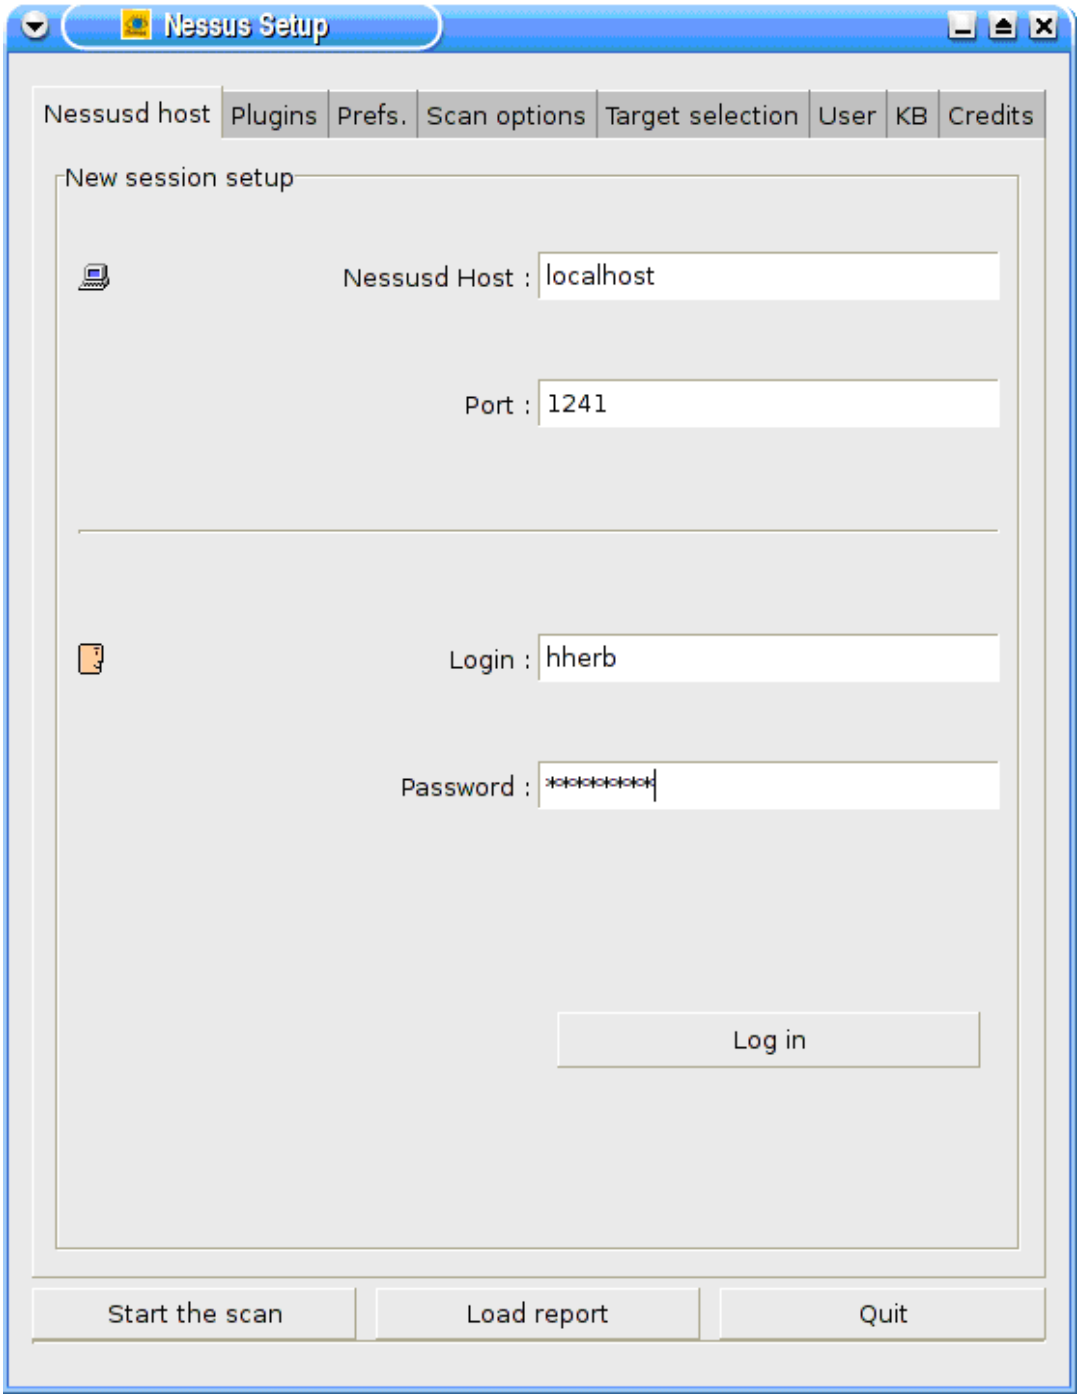

4 If you start the Nessus program for the first time, it will ask you whether you accept the server certificate. Say yes, because it is the one you just created before with 'nessus-mkcert'.

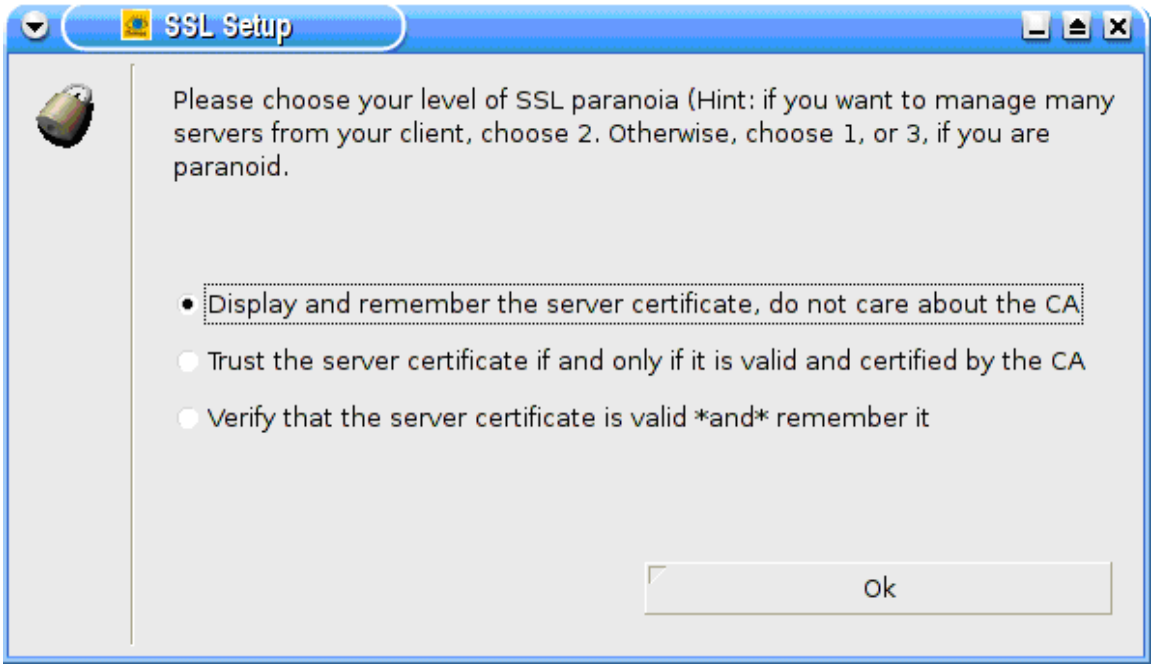
SSL Setup
Please choose your level of SSL paranoia (Hint: if you want to manage many servers from your client, choose 2. Otherwise, choose 1, or 3, if you are paranoid.
Display and remember the server certificate, do not care about the CA
Trust the server certificate if and only if it is valid and certified by the CA
Verify that the server certificate is valid *and* remember it
Ok

5 If login is successful you will presented with the certificate for visual verification. In the scenario here it is safe again to simply click OK.

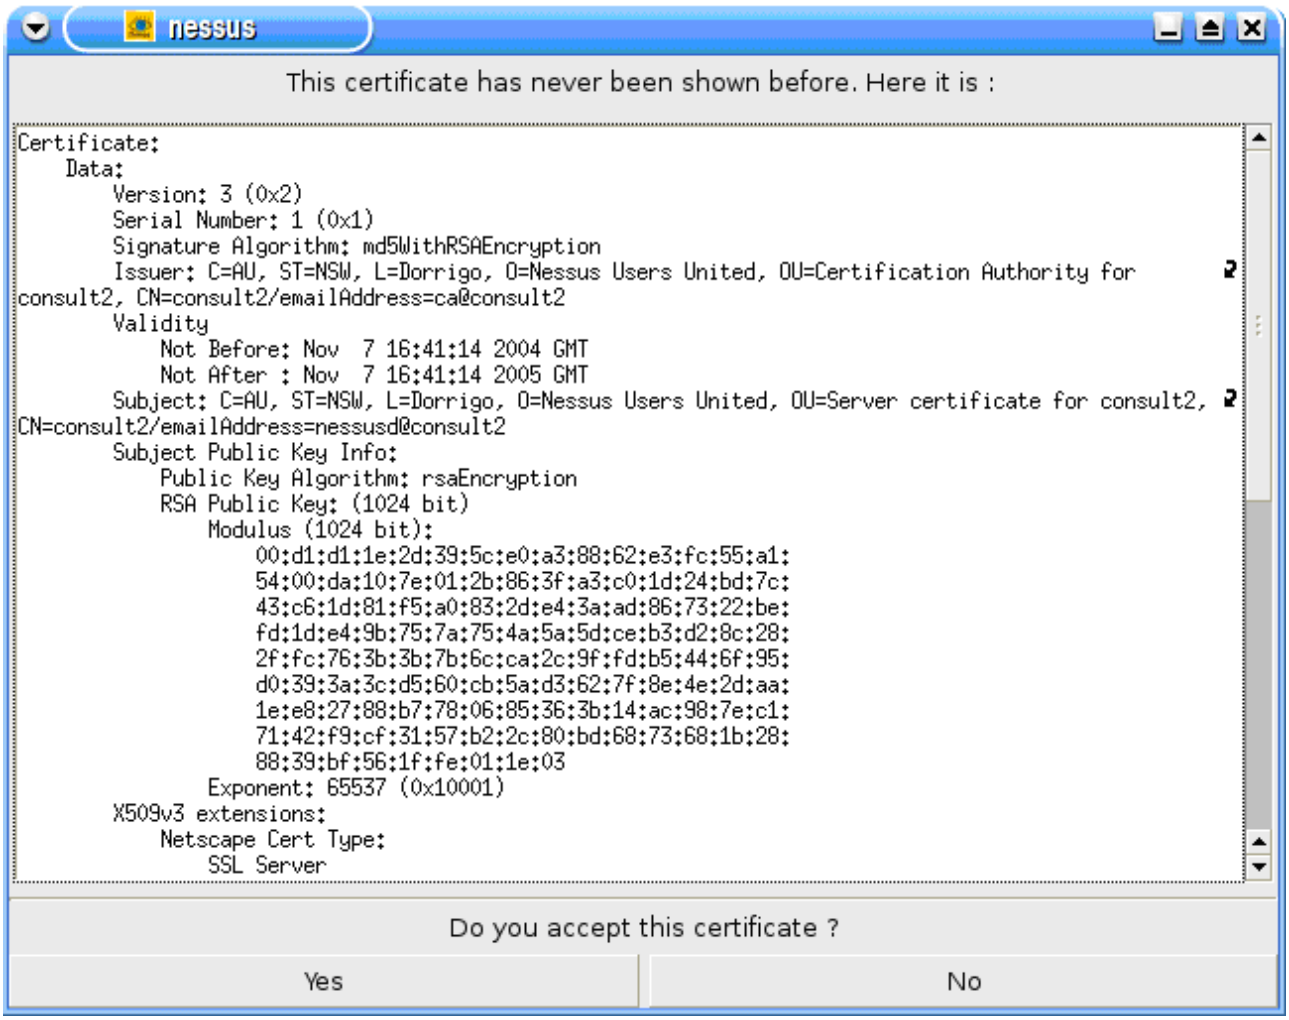

A warning will probably pop up telling you that dangerous features have been disabled (those which might crash a victim during scanning). Accept this for now.

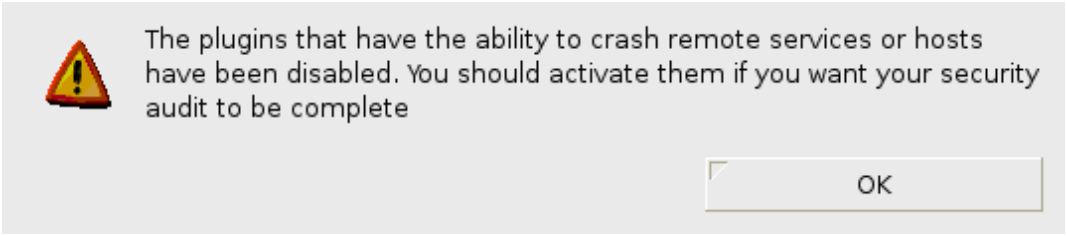
The plugins that have the ability to crash remote services or hosts have been disabled. You should activate them if you want your security audit to be complete
OK

6 Time to quickly review the configuration.

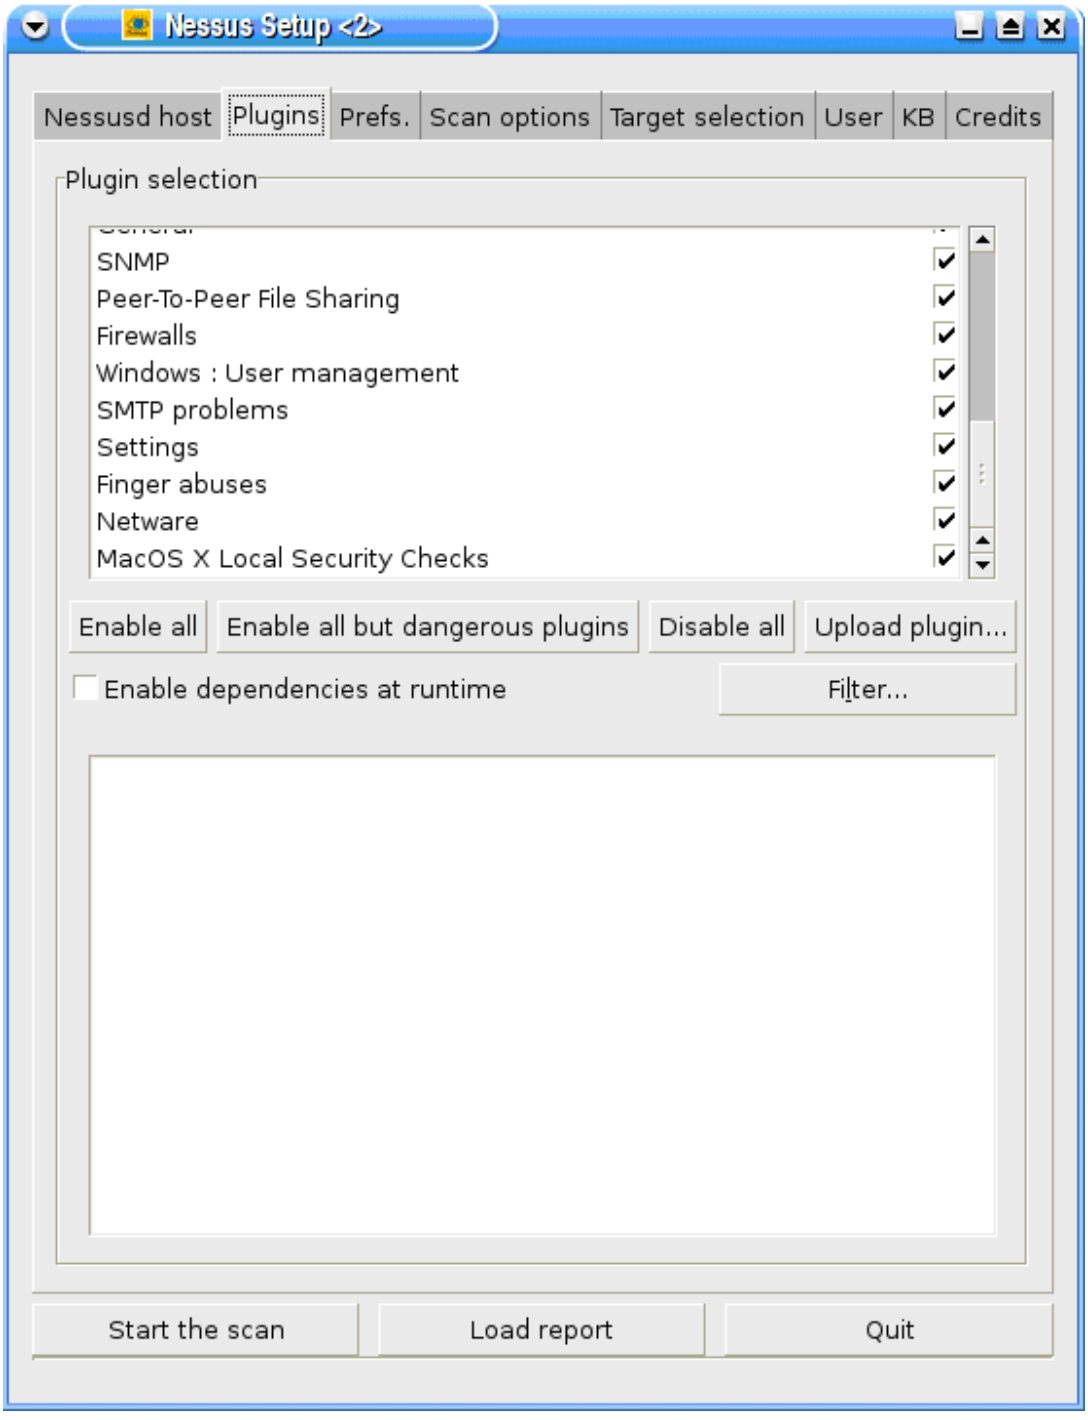

7 In the plugin section, simply enable ‘all but dangerous plugins’ for now. For the first scan, you can leave all other configuration options at their default settings.

8 The last thing to do before our first scan is to select the target. Instead of the depicted ‘my.firewalls.ip’ you would enter the public IP address or hostname of your firewall.

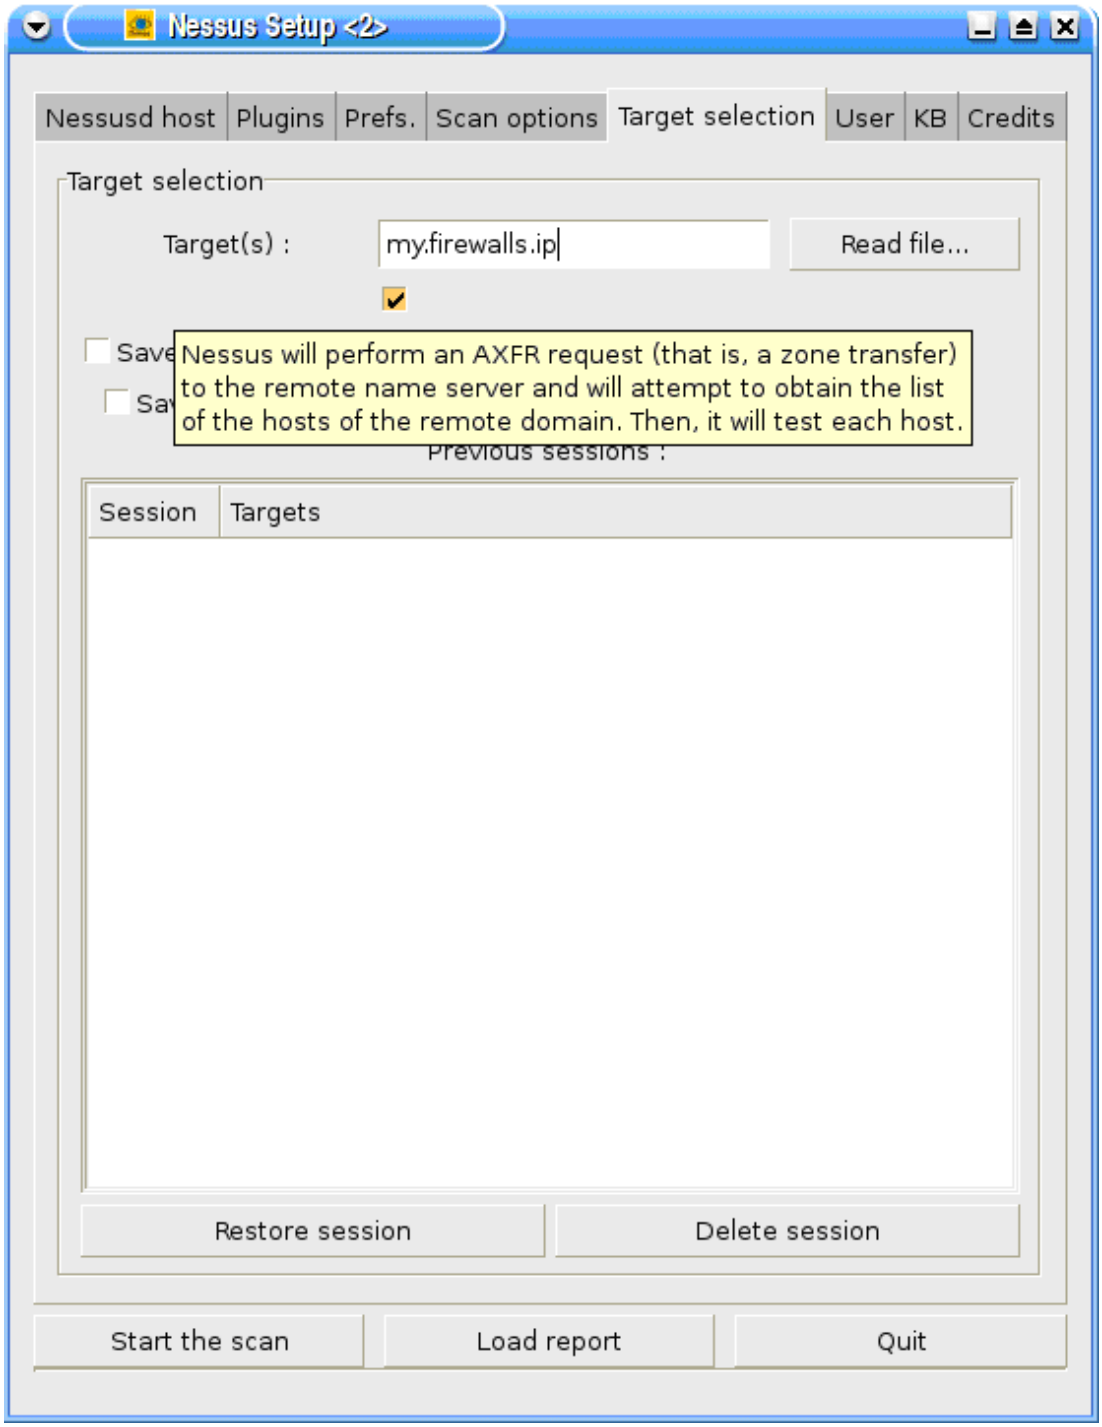

The Nessus project web site has more information.

9 Now, all that is left to do is to click on the ‘Start the scan’ button at the bottom of the Nessus dialog box. It may take anything from several minutes to several hours. A progress bar will indicate progress. Once finished, a very detailed test report will be displayed.

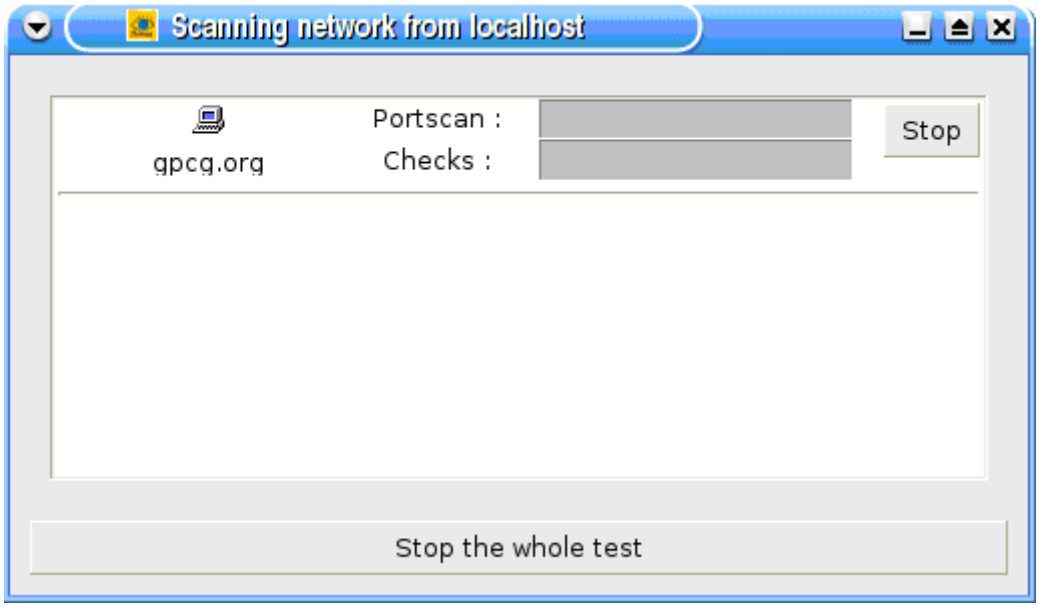

STEP 12: Firewall checklist – after installation

After installation and configuration of your firewall, but before you connect your private network to the Internet via your firewall, please go through this checklist.

If there is even one question you cannot answer with yes, reconsider your options before connecting to the Internet.

1 Is the firewall the only device with a network interface that can connect to the Internet?
- Make sure there are no modems connected to computers within the private network.
- Make sure there is no hub or switch bypassing the firewall.

2 Have you updated the firewall to the newest firmware/software version?
- If not, consult your manual how to do this, and check on the Internet for availability of such updates. Running old software with possibly known vulnerabilities is dangerous.

3 Have all ports on the firewall been closed, unless there is a demonstrated need for keeping a specific port open, and has a risk assessment been performed?
- If not, please read about firewall configuration in Step 9 first.

4 Has the firewall been tested?
- If not, read first about firewall auditing in Step 10.
- Then follow our step by step auditing guideline in Step 11 (if you cannot get a professional to do it).

5 Has the firewall passed the tests without any warnings/failures?
- If not, you must fix the detected problems first and retest your firewall until the test report comes back clear.

6 If you could answer all questions so far with a ‘yes’, proceed. Otherwise, go back to question 1.
- If you don't take this seriously, you will most likely suffer an intrusion and subsequently be in breach of the Privacy Act.

7 Now you may connect your private network to the firewall.
- Can you access all Internet services you intended to use?
  - can you browse the web?
  - can you send and receive email?
  - can you download pathology results etc?
- If everything works, you must perform one final firewall audit on your live configuration because computers active behind the firewall may show risks that weren’t obvious before.

FURTHER INFORMATION

Virtual Private Network

Sometimes it may be useful to extend your private network outside your practice building – for example, to connect to a branch surgery, to access your practice network from home or while travelling, or from the local hospital.

In most cases, it is not be possible to extend your Ethernet cables to that other location. Sometimes, when there is line of sight, a wireless link might be feasible, but usually distances and geography do not allow this.

However, as long as both locations can connect to the Internet, they can also connect to each other. Since the Internet is a public place, you have to take some precautions to keep your private network traffic private. To do this, you have to create a Virtual Private Network (VPN), or a secure tunnel through the Internet.

On a properly configured VPN, you don't have to distinguish between your local and the remote network. For practical reasons, especially regarding privacy, they can be considered as a single cabled network.

This is achieved by strongly encrypting all network packets flowing between the two private locations. In the ideal case, this is done transparently by a ‘VPN router’, which is a device that automatically encrypts all traffic flowing from local to remote, and automatically decrypts all traffic flowing from remote to local.

Such device guarantees that all traffic flowing through the Internet between the two VPN endpoints is kept strictly confidential, regardless of what software is used on these computers (e.g. email, word processor etc.).

Figure 6: A VPN router transparently encrypts and decrypts network traffic

There are many different protocols available for VPN implementation. However, for the time being, only consider the international Internet standard protocol, called IPSec.

Microsoft initially had its own VPN protocol, called PPtP. However, it is considered insecure and should never be used in a medical practice setting. For this reason, do not use any of the VPN routers available which are only capable of PPtP instead of IPSec.

The most common scenario for a VPN might be linking two branch surgeries together. As depicted in Figure 7, this can be achieved easily through two properly configured and compatible VPN routers. Nothing needs to be changed software-wise on any of the computers in either practice.

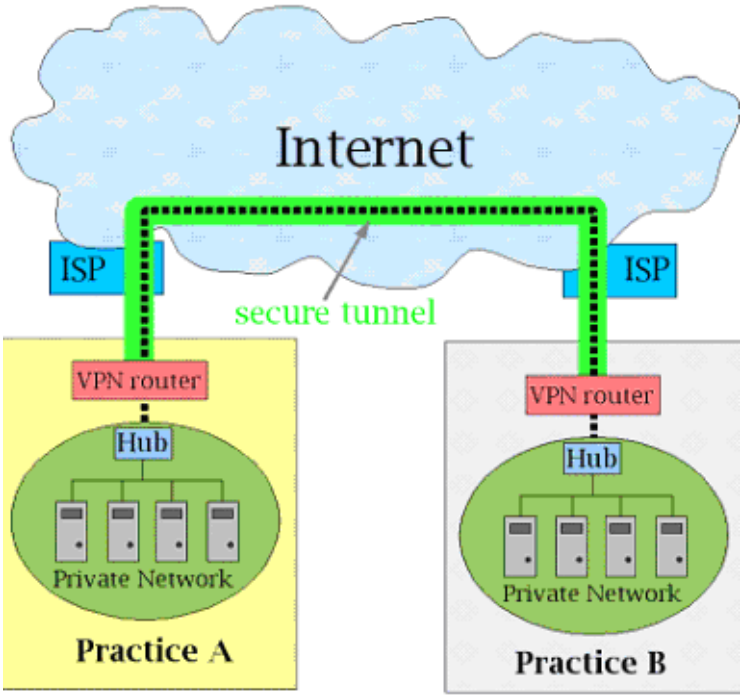

Figure 7: Connecting two practices via VPN

Another common scenario is connecting to the practice from home or while travelling, using a notebook and a dial-up connection for example. Unless a properly configured VPN router is also carried along, some software will need to be installed in this case. You can save yourself time and effort if your VPN router is compatible with the IPSec standard and does not depend on (costly and possibly in the future unsupported) proprietary client software.

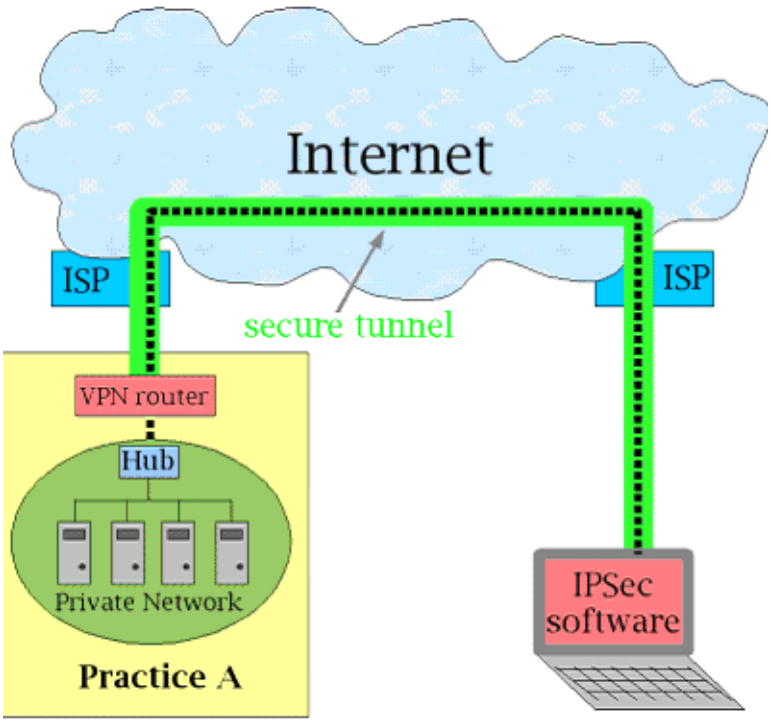

Figure 8: Connecting a laptop to the practice via VPN

Failover/load balancing

Failover

The failover principle is to have multiple Internet service providers, and let your gateway device handle the connections for you automatically, depending on needs and availability of service.

The problem

We are becoming increasingly dependent on the Internet. Email is becoming the mainstream communications medium, not only replacing traditional postal mail but also phone communications to some degree. Financial transactions (online banking and shopping) is gaining importance, as is online information access of any kind.

Especially in medical practice, we expect online services (like pathology result downloads) to be accessible around the clock. However, often this depends on access to developing broadband infrastructure and types of online services offered by the many vendors involved.

In cases where you absolutely rely on broadband connectivity (as in hosting our medical records with an Application Service Provider (ASP) or via VPN at a branch surgery) you cannot afford to depend on a single unreliable service. Even with uptime of 99.9%, you still have one working day per year when the service will not be available from this provider, and virtually no Australian service provider can even go as far as providing a 99.9% uptime guarantee.

The solution

The solution is not putting all your eggs into a single basket. Choose one Internet access service as your primary service, but always have at least one more service that uses different infrastructure as backup.

Examples of services using different infrastructures:

- (A)DSL.
- phone line dial up – strictly speaking it uses the same infrastructure as (A)DSL (namely your phone line), but the most frequent cause of failure of the (A)DSL system will not disrupt phone services on the same line at the same time.
- cable Internet.
- two way Satellite Internet (VSAT).
- wireless Internet.

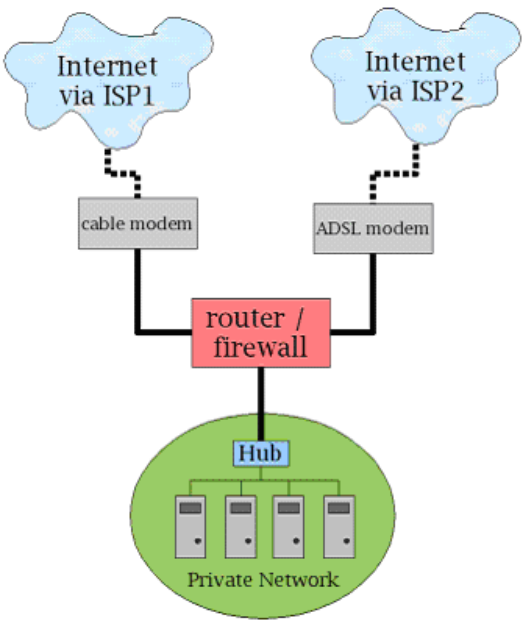

Figure 9: accessing two different ISPs via a single router device with 2 WAN ports

Unfortunately, there is no rule regarding which technology is the most reliant at present in Australia. It depends on a variety of technological and vendor specific factors that vary from location to location.

If your location allows it, the combination of cable and (A)DSL would usually be the choice with the best performance and lowest cost. In some locations, a wireless Internet provider might be an even better substitute for either cable or (A)DSL.

Dual WAN ports

Some Internet gateway appliances (embedded router/firewall devices) allow more than one WAN (in this case, ISP) connection. They can be set up to automatically switch over to the secondary provider if the primary provider is not available.

For example – you have an (A)DSL provider and a cable provider. You connect WAN port 1 from your gateway with your cable modem (since your cable service is faster and cheaper to use), and WAN port 2 of your gateway to your (A)DSL modem. You configure your gateway so that cable is used whenever available, with (A)DSL as a backup that is used automatically when cable is unavailable, until cable becomes available again.

Load balancing

If you have two independent Internet access methods, you may as well make the best of it. Why not boost your performance by using both, instead of just using the secondary connection if the primary connection fails?

Imagine somebody in your practice is downloading a huge radiology image, and your downloading software is not smart enough to share bandwidth fairly. Everybody else would be almost locked out of Internet use until the download finishes.

Load balancing can be done in principle with a single connection on a packet-per-IP basis, but it is far more efficient and predictable if you have more than one physical connection to the Internet.

Load balancing in this case means that your Internet gateway (your firewall/router) takes care of everybody getting a fair share of the available bandwidth.

If your router has load balancing features, it will allow you to route the network traffic through either connection depending on criteria you can select in the router’s configuration:

- if load exceeds a certain percentage.
- on a particular day of the week or time of the day.
- depending on network traffic volume (there is a potential cost saving if you can avoid traffic limits imposed by your contract).

GLOSSARY

Access – The ability to use computer information in some manner. Specific access can be granted to each individual user.

Application services - Services that leverage bandwidth to deliver increased functionality and value to subscribers.

ASP – Application Service Provider. A third party entity that manages and distributes software-based services and solutions to customers across a wide area network from a central data centre.

Biometrics - The technique of studying physical characteristics of a person, such as fingerprints, hand geometry, eye structure or voice pattern.

Defence Signals Directorate (DSD) - National authority for signals intelligence and information security.

DMZ – Demilitarised Zone **-** A firewall configuration for securing local area networks (see LAN)

DNS – The Domain Name Server is a system that translates domain names into IP addresses.

Encryption – A procedure used to convert plaintext into ciphertext in order to prevent any but the intended recipient from reading that data. There are many times of data encryption, and they are the basis of network security.

Failover – A backup operation, which will automatically switch to a standby system if and when a primary system fails. These standby systems may include databases, server, and networks.

Firewall – Systems used to prevent unauthorised access. The firewall may be hardware, software or both.

Firmware - Software stored in the computers read only memory (ROM) and cannot be changed.

FTP – File Transfer Protocol. A protocol which allows a user on one host to access, and transfer files to and from, another host over a network.

Hardware - The physical equipment of a computer system, including the monitor, keyboard, central processing unit, and storage devices.

HTTP - Hyper Text Transfer Protocol. The WWW protocol that performs the request and retrieve functions of a server. Commonly seen as the first part of a website address.

Internet - An interconnected system of networks that connects computers around the world via the TCP/IP protocol.

IP address - Internet Protocol Address. A series of four numbers between one and 3 digits in length, numbers separated by periods. It is used to identify a computer connected to the Internet. For example, 212.6.125.76 is an IP address.

IPsec – Short for IP Security. A set of protocols to support secure exchange of packets at the IP level.

IP spoofing – Using a legitimate IP address or packet to gain unauthorised access to a computer.

ISP - Internet Service Provider. A company that provides an Internet connection.

LAN - A local area network (LAN) is a computer network covering a local area, for example a home, general practice or small group of buildings such as a health centre. The topology of a network dictates its physical structure.

Load Balancing – The distribution of processing and communications activities evenly across a computer network to avoid a single device becoming overwhelmed.

Malware - Malicious code, in the form of viruses, worms, and Trojan Horses.

Modem - This is short for modulator-demodulator devices. Modems allow computers to transmit information to one another via an ordinary telephone line.

Network – Two or more computers linked together.

Network gateway - An inter-networking system that joins two networks together. A network gateway can be implemented completely in software, completely in hardware, or as a combination of the two.

Network interface – A boundary across which two independent systems meet and communicate with each other.

Packet - A bundle of data organized for transmission, containing control information (destination, length, origin, etc.) the data itself and error detection and correction bits.

Patches - Upgrades for software supplied by manufacturers such as Microsoft usually over the Internet.

Port – In a communications network, a port is a point at which signals can enter or leave the network en route to or from another network.

Protocol - A formal description of message formats and the rules two computers must follow to exchange those messages.

Router - An electronic device that connects two or more networks and routes incoming data packets to the appropriate network.

Rulebase – Component of logic system that specifies the meanings of the well-formed expressions of the logical language.

Ruleset - A rule set contains an ordered group of configured rules which are sequentially tested or applied.

Security – Techniques for ensuring that any data stored on a PC can not be read or compromised by individuals who don't have authorisation. Most PC security includes passwords and data encryption.

SMPT - Simple Mail Transfer Protocol is the protocol used to transfer e-mail.

Software - A program or set of instructions that controls the operation of a computer. Distinguished from the actual hardware of the computer.

Spam mail - The Internet version of junk mail. Spamming is sending the same message to a large number of mailing lists or newsgroups usually to advertise something.

Spyware - Any software that covertly gathers user information through the user's Internet connection without his or her knowledge, usually for advertising purposes.

Stateful Inspection – Also know as dynamic packet filtering, stateful inspection is a firewall architecture that provides enhanced security by keeping track of and examining both incoming and outgoing packets.

TCP/IP – Transmission Control Protocol/Internet Protocol.

Trojan Horse - a sort of virus although it does not replicate itself - it is something which is hidden in a file and, when activated, it does terrible damage ... like a Friday 13th thingy - obviously named after the Trojan Horse.

Virus - In computer security technology, a virus is a self-replicating program that spreads by inserting copies of itself into other executable code or documents.

VOIP – Voice Over Internet Protocol. Using the internet protocol to carry voice data. Depending on the scenario, VOIP can facilitate cheap or even free phone calls.

VPN – Virtual Private Network. A private data network that makes use of the public telecommunications infrastructure, maintaining privacy through the use of a tunnelling protocol and security procedures.

VPN IPSEC - Internet Protocol Security is an internationally recognised VPN protocol suite developed by the IETF (Internet Engineering Task Force).

Worm - A computer worm is a self-replicating computer program, similar to a computer virus. A virus attaches itself to, and becomes part of, another executable program; however, a worm is self-contained and does not need to be part of another program to propagate itself

General Practice Computing Group
C/- Royal Australian College of General Practitioners
1 Palmerston Crescent
South Melbourne, Vic 3205
Tel: (03) 8699 0414

www.gpcg.org.au